arXiv:1401.0647v2 [cs.SC] 24 Apr 2014

Cylindrical Algebraic Sub-Decompositions

D. J. Wilson, R. J. Bradford, J. H. Davenport and M. England

Abstract. Cylindrical algebraic decompositions (CADs) are a key tool in real algebraic geometry, used primarily for eliminating quantifiers over the reals and studying semi-algebraic sets. In this paper we introduce cylindrical algebraic sub-decompositions (sub-CADs), which are subsets of CADs containing all the information needed to specify a solution for a given problem.

We define two new types of sub-CAD: variety sub-CADs which are those cells in a CAD lying on a designated variety; and layered sub-CADs which have only those cells of dimension higher than a specified value. We present algorithms to produce these and describe how the two approaches may be combined with each other and the recent theory of truth-table invariant CAD.

We give a complexity analysis showing that these techniques can offer substantial theoretical savings, which is supported by experimentation using an implementation in Maple.

Keywords. Cylindrical Algebraic Decomposition, Real Algebraic Geometry, Equational Constraints, Symbolic Computation, Computer Algebra.
MSC Code: 68W30 (Symbolic Computation and Algebraic Computation).

1. Introduction

1.1. Motivation

A **cylindrical algebraic decomposition** (CAD) is a decomposition of $\mathbb{R}^n$ into cells arranged cylindrically (meaning the projections of any pair of cells onto the first k coordinates are either equal or disjoint) each of which is a semi-algebraic set (and so may be described by polynomial relations). They are traditionally produced sign-invariant with respect to a list of polynomials meaning each polynomial has constant sign on each cell. CAD was introduced by Collins in [17], and has become a key tool in real algebraic geometry for studying semi-algebraic sets and eliminating quantifiers over the reals. Other applications include robot motion planning [38], parametric optimisation [27], epidemic modelling [12], theorem proving [36] and programming with complex functions [21].

CAD usually produces far more information than required to solve the underlying problem. Often a problem will be represented by a formula and we thus require a CAD such that the formula has constant truth value on each cell. This can be achieved by building a CAD sign-invariant for the polynomials in the formula, but that may introduce cell divisions not relevant to the formula itself. Many techniques have been developed to try and mitigate this, some of which we discuss later. However, even then algorithms may produce thousands of superfluous cells which are not part of the solution set. The key focus of this paper is the development of methods to return a subset of a CAD sufficient to solve a given problem. We show that such subsets can often be identified from the structure of the problem, motivating the following new definitions.

Definition 1.
Let $\mathcal{D}$ be a CAD of $\mathbb{R}^n$ (represented as a set of cells). Then a subset $\mathcal{E} \subseteq \mathcal{D}$ is a **cylindrical algebraic sub-decomposition (sub-CAD)**.

Let $F \subset \mathbb{Q}[x_1, \ldots, x_n]$. If $\mathcal{D}$ is a sign-invariant CAD for F then $\mathcal{E}$ is a **sign-invariant sub-CAD**. We define sub-CADs with other invariance properties in an analogous manner, such as **truth-invariance** for a Tarski formula $\varphi(x_1, \ldots, x_n)$. If $\mathcal{D}$ is a truth-invariant CAD for φ and $\mathcal{E}$ contains all cells of where φ is satisfied then we say that $\mathcal{E}$ is a **φ-sufficient sub-CAD**.

As an example of when sub-CADs may be applicable, consider quantifier elimination, the original motivation for CAD. Given a quantified formula φ we want to derive an equivalent quantifier-free formula. For a formula over the reals this is achieved by constructing a sign-invariant CAD for the polynomials in φ and testing the truth of φ at a sample point of each cell. This is sufficient to draw a conclusion for the whole cell due to sign-invariance and thus an equivalent quantifier free formula can be created from the algebraic description of the cells on which φ is true. Such an application makes no use of the cells on which φ is false and so a φ-sufficient sub-CAD is appropriate.

Of course, for a given problem we would like the smallest possible φ-sufficient sub-CAD. It is not usually possible to pre-identify this, but we have developed techniques which restrict the output of the CAD algorithm to provide sub-CADs sufficient for certain general classes of problems. These will offer savings on any subsequent computations on the cells (such as evaluating polynomials or formulae) and in some cases also offer substantial savings in the CAD construction itself. We will introduce these techniques and demonstrate how they can be combined with each other and additional existing CAD theory, but first remind the reader of the necessary background theory.

1.2. Background to CAD

Collins' original algorithm is described in [1]. While there have been many improvements and refinements to this algorithm the structure has remained largely the same. In the first phase, **projection**, a **projection operator** is repeatedly applied to a set of polynomials, each time producing another set in one fewer variables. Together these sets contain the **projection polynomials**. These are then used in the second phase, **lifting**, to build the CAD incrementally. First $\mathbb{R}$ is decomposed into cells: points corresponding to the real roots of the univariate polynomials, and the open intervals defined by them. Then $\mathbb{R}^2$ is decomposed by repeating this process over each cell using the bivariate polynomials (evaluated at a sample point). The output for each cell consists of **sections** (where a polynomial vanishes) and **sectors** (the regions between). Together these form a **stack** over the cell, and taking the union of these stacks gives the CAD of $\mathbb{R}^2$. This is repeated until a CAD of $\mathbb{R}^n$ is produced. To conclude that the CAD is sign-invariant we need delineability. A polynomial is **delineable** over a cell if the portion of its zero set over that cell consists of disjoint sections. Then a set of polynomials is **delineable** over a cell if each is delineable and the sections of different polynomials over the cell are either identical or disjoint. The projection operator used must ensure that over each cell of a sign-invariant CAD for the projection polynomials in r variables, the polynomials in $r + 1$ variables are delineable.

All cells include a cell index and a sample point. The index is an n-tuple of positive integers that corresponds to the location of the cell relative to the rest of the CAD. Cells are numbered in each stack during the lifting stage (from most negative to most positive), with sectors having odd numbers and sections having even numbers. Therefore the dimension of a given cell can be easily determined from its index: simply the number of odd indices in the n-tuple. Our algorithms in this paper will produce sub-CADs that are **index-consistent**, meaning a cell in a sub-CAD will have the same index as it would in the full CAD. Further, we will assume that cells are stored lexicographically by index.

Important developments to CAD include: refinements to the projection operator [28, 33, 7], reducing the number of projection polynomials and hence cells; partial CAD [18], where the structure of the input formula is used to simplify the lifting stage; the theories of equational constraints and truth-table invariance [34, 4] where the presence of equalities in the input further refines the projection operator; the use of certified numerics in the lifting phase [41, 29]; and CAD via triangular decomposition [16] which constructs a decomposition of complex space and refines this to a CAD.

Constructing a CAD is doubly exponential in the number of variables [22]. While none of the improvements described above (or introduced in this paper) circumvent this they do make a great impact on the practicality of using CAD. Note that CAD can depend heavily on the variable ordering

used (from linear to doubly-exponential [11]). In this paper we work with polynomials in $\mathbb{Q}[\boldsymbol{x}]$ with the variables $\mathbf{x} = x_1, \ldots, x_n$ in ascending order (so we first project with respect to x_n and continue until we reach univariate polynomials in x_1). The **main variable** of a polynomial (mvar) is the greatest variable present with respect to the ordering. Heuristics to assist with selecting a variable ordering (and other choices) were discussed in [23, 6] and are equally applicable to sub-CADs.

1.3. New Contributions

In Section 2 we present new algorithms to produce sub-CADs, as well as surveying the literature to identify other examples of sub-CADs. To the best of our knowledge the concept of a sub-CAD has never been formalised and unified before.

We start in Section 2.1 by defining a **Variety sub-CAD (V-sub-CAD)**. This idea combines the ideas of: equational constraints [34], where the presence of an equation implied by the input formula improves the projection operator; and partial CAD [18], where the logical structure of the input allows one to truncate the lifting process when the truth value can already be ascertained. We observe that if the input formula contains an equational constraint then all valid cells must lie on the variety it defines and hence it is unnecessary to produce cells not on this variety.

In Section 2.2 we define a **Layered sub-CAD (L-sub-CAD)** as the cells in a CAD of a specific dimension or higher. It has been noted previously that a problem involving only strict inequalities would require only the cells in a CAD of full-dimension [31, 40]. We generalise this idea and explain when it may be of use, for example to solve problems whose solution sets are of known dimension or in applications like robot motion planning where only cells of certain dimensions are of use.

These new ideas improve the practicality of using CAD and their effect can be increased by combining them, as discussed in Section 3.1. For example, consider formulae of the form $f = 0 \wedge \varphi$ where φ involves only strict inequalities. Then a **Layered Variety Sub-CAD (LV-sub-CAD)** can provide the cells of full dimension on the variety and thus the generic families of solutions.

These new ideas may also be combined with many existing aspects of CAD theory. It is of course sensible to combine the restricted output of a variety CAD with the theory of reduced projection with respect to an equational constraint. However, it is also possible to combine with the projection operator for **truth-table invariant CAD (TTICAD)** recently presented in [4]. A TTICAD is one for which each cell is truth invariant for a list of formulae, utilising equational constraints in the individual formulae to reduce the number of projection polynomials. We discuss when and how truth-table-invariant sub-CADs can be produced in Section 3.2. In Section 5.2 we examine a problem where a LV-sub-TTICAD can be used to identify almost all the solutions (the set of missing solutions has measure zero). This approach produces 88% fewer cells than using TTICAD alone and takes seconds rather than minutes (while trying to tackle the problem with a traditional CAD is infeasible). This and two other case studies demonstrating the benefit of the new algorithms are presented in Section 5.

In Section 4 we give a complexity analysis of certain sub-CADs. Although none of the new theory allows us to avoid the doubly exponential nature of CAD they do allow for improved asymptotic bounds. The improvement is a drop in the constant term of the double exponent, and we note that $2^{2^n} \neq O(2^{2^{n-1}})$ so such savings can have a substantial effect. This is reflected by experimental results in Section 5 where substantial increases in efficiency due to sub-CAD technology are demonstrated.

2. Sub-CADs

We aim to return only those cells necessary to solve the problem at hand: a ϕ-sufficient sub-CAD. However, trying to identify the minimal ϕ-sufficient sub-CAD for a problem would mean essentially solving the problem itself and so we instead explain how to identify sets of valid or invalid cells during the lifting stage at minimal cost. Indeed, the two new approaches to sub-CAD we present require only simple checks on cell-dimensions (easily obtained via the cell-index). We present theory and algorithms for variety and layered sub-CADs in Sections 2.1 and 2.2, and then in Section 2.3 we put our work into context by surveying the CAD literature for relevance to sub-CADs.

2.1. Variety sub-CADs

Recall the definition of an equational constraint.

Definition 2.
Let φ be a Tarski formula. An **equational constraint** is an equation, $f = 0$, logically implied by φ.

Equational constraints may be given explicitly (as in $f = 0 \wedge \phi$), or implicitly (as $f_1 f_2 = 0$ is in $(f_1 = 0 \wedge \phi_1) \vee (f_2 = 0 \wedge \phi_2)$). The presence of an equational constraint can be utilised in the first projection stage by refining the projection operator [34] and also in the final lifting stage by reducing the amount of polynomials used to construct the stacks [25]. If more than one equational constraint is present then further savings may be possible [35, 13]. We restrict ourselves to a single equational constraint, and if multiple equational constraints are present we assume that one has been designated.

Definition 3.
Let φ be a Tarski formula with equational constraint $f = 0$. A truth-invariant sub-CAD for φ consisting only of cells lying in the variety defined by $f = 0$ is a **Variety Sub-CAD (V-sub-CAD)**.

Partial CAD [18] describes how the logical structure of the input formula is used to truncate lifting when possible in CAD. For example, if the truth of an expression on a cell c can already be determined then there is no need to lift over it. Algorithm 1 combines the ideas of utilising equational constraints and partial lifting to build variety sub-CADs in the case where all factors of the equational constraint have the main variable of the system.

In Algorithm 1 A is a square-free basis for the polynomials defining φ and E the subset of those defining f. `ProjOp` refers to an algorithm implementing a suitable CAD projection operator. Then `CADAlgo` and `GenerateStack` respectively implement compatible algorithms for CAD construction, and stack generation over a cell with respect to the sign of given polynomials. By compatible we mean using the same projection operator and checking for any necessary conditions of its use. This is required as some CAD algorithms may return FAIL if the input does not satisfy certain conditions. These are checked for during stack construction and usually referred to as the input being **well-oriented** (see for example [33]). In these cases Algorithm 1 must also return FAIL. A sensible choice of projection operator is one which utilises the equational constraint to minimise the number of projection polynomials, such as $P_E(A)$ from [34]. We verify the correctness of Algorithm 1 for this choice.

Theorem 1. *When the sub-algorithms are chosen to implement McCallum's algorithm to produce CADs with respect to an equational constraint* [34]*, then Algorithm 1 satisfies its specification, with the outputted sub-CAD consisting of cells on which the input formula has constant truth value.*

Proof. The algorithm in [34] applies a projection operator $P_E(A)$ and then incrementally constructs CADs of increasing real dimension, checking for well-orientedness when building each stack. In [34] the authors proved that the CAD returned was truth-invariant for the equational constraint and sign-invariant for any other polynomials involved on cells where the equational constraint was satisfied.

The first difference in Algorithm 1 is in step 8 where the final lift is performed with respect to E rather than A. In fact, this improvement follows directly from Theorem 2.2 in [34], although it was not realised until [5]. This reduces the size of the output, but not its invariance structure or correctness.

Next, in the final loop, only some of the cells generated in the final lift are included in the output. The cells in question are a subset of what would have been produced otherwise and thus certainly a sub-CAD with the same invariance property. It remains to prove that they are a variety sub-CAD (in which case we can conclude φ has constant truth value on each cell).

The selected cells are those with even index (the sections). If the polynomial f is not identically zero over a cell in $\mathcal{D}$ then these must together define its variety. If any of the polynomials in E were nullified then part of the variety may be in the sectors, but in this case the input would have failed the well-orientedness condition in [34] and thus Algorithm 1 would return FAIL. □

Remark 1. As the operator from [34] can return FAIL in situations where others do not, we now consider how Algorithm 1 may be adapted to use alternative CAD projection operators.

Algorithm 1: `VarietySubCAD`$(\varphi, f, \mathbf{x})$: Algorithm to produce variety sub-CADs.

Input : A formula φ, a declared equational constraint $f = 0$ from φ and variables $\boldsymbol{x} = x_1, \ldots, x_n$. φ is in $\boldsymbol{x}$ and all factors of f have main variable x_n.
Output: A (truth-invariant) variety sub-CAD of $\mathbb{R}^n$ for (φ, f), or FAIL.

```
1  Extract from φ the set of polynomials A and from f the subset E ⊂ A;
2  P ← output from applying ProjOp to (A,E) once ;          // First projection stage
3  D' ← CADAlgo(P, [x_1, ..., x_{n-1}]) ;                   // Computation of a CAD of R^{n-1}
4  if D' = FAIL then
5      return FAIL ;                                        // P is not well oriented
6  D ← [];
7  for c ∈ D' do
8      S ← GenerateStack(E, c);                             // Final lifting stage
9      if S = FAIL then
10         return FAIL ;                                    // Input is not well oriented
11     if |S| > 1 then
12         for i = 1 ... (|S| − 1)/2 do
13             D.append(S[2i]) ;                            // Cells with even index included
14 return D;
```

First, if a polynomial in E is nullified on a cell of $\mathcal{D}$ then [34] returns FAIL while McCallum's operator to produce sign-invariant CADs in [33] is still applicable (because then the nullification is in the final lift where only sign-invariance and not order invariance is required). However, we cannot simply apply Algorithm 1 with the alternative projection operator as it will now be the case that some of the variety is contained in the sectors over the cell in question. In this case we would need the `GenerateStack` algorithm to check for nullification of E, and then if it occurs have Algorithm 1 include all cells from that stack in the output.

Second, if some other polynomial is nullified causing failure then the original algorithm of Collins (or its improvement by Hong [28]) is still applicable. As with the previous case we must still check for nullification over a cell in the final lift, including the whole stack when nullification occurs.

In these cases there would still be output savings from building a variety sub-CAD since the inclusion of the full stack only needs to happen over those cells where nullification occurs.

We demonstrate the savings in output size offered with a simple example.

Example 1.
Assume variable ordering $x \prec y$ and define the polynomials

$$f := x^2 + y^2 - 1, \qquad g := x$$

which are graphed with solid curves in each of the images in Figure 1. This first of these images visualises the simplest sign-invariant CAD for the polynomials, which has 23 cells (each indicated by a solid box). If the box is at the intersection of two curves (including the dotted lines) then the cell it represents is a point. Otherwise, if the box is on a curve then the cell represented is that portion of the curve and if the box is not on a curve then the cell represented is that portion of the plane.

Suppose these polynomials originated from a problem involving the formula $\varphi_1 := f = 0 \wedge g < 0$. Then 3 of the cells describe the solution (those on the circle to the left of the y-axis). Using Algorithm 1 a V-sub-CAD would be returned with all 8 of the cells on the circle. The other cells have been recognised as not being on the variety to which all solutions belong and hence have been discarded. The image on the right in Figure 1 identifies the cells returned in the sub-CAD with solid boxes.

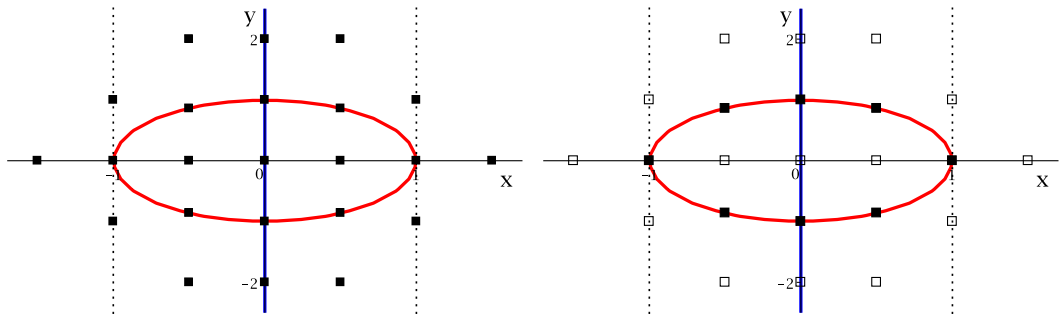

FIGURE 1. Images representing the CADs described in Example 1. The solid boxes represent cells in $\mathbb{R}^2$, (with the empty boxes cells that were constructed but discarded).

The cell counts in this example were modest, but we will see later in Section 5 that for more complicated examples the savings offered by using a V-sub-CAD can be substantial. Using Algorithm 1 clearly reduces the output size of CAD but it will save little CAD computation time since most of the work required to define the discarded cells (such as root isolation) had to be performed. These computed but discarded cells are shown as empty boxes in Figure 1. There would, however, be savings in computation time for any further work, such as only having to evaluate polynomials on 8 cells instead of 23 in order to describe the solutions.

If a formula has an equational constraint with factors not in the main variable of the system (parts of the corresponding variety have lower dimension) then building a V-sub-CAD would mean even more potential cell savings, but also savings in computation time since some stacks would never be built at all. Suppose each factor of the equational constraint has main variable x_k where $k < n$. To adapt Algorithm 1 we must perform the restriction when lifting to a CAD of $\mathbb{R}^k$, instead of $\mathbb{R}^n$. Thus we would first build a CAD of $\mathbb{R}_{k-1}$, then perform the restricted lifting to a sub-CAD of $\mathbb{R}^k$, and continue lifting to a sub-CAD of $\mathbb{R}^n$. However, verifying this approach is a little more subtle. We cannot follow Theorem 1 and use McCallum's reduced projection at the start (as the set E is empty [34]) but if we were to use it to build a sub-CAD of $\mathbb{R}^k$ then we would need to take care in how we lift over it (ensuring the projection polynomials above are delineable). We see two possibilities:

(a) Use the tools of McCallum's sign-invariant algorithm from [33] throughout. In particular, we must perform the restricted lifting with respect to the full set of projection polynomials of main variable x_k rather than just those defining the equational constraint. This is because to continue lifting with respect to projection polynomials provided by McCallum's operators we need to conclude that the sub-CAD of $\mathbb{R}^k$ is order invariant on the cells, rather than just sign-invariant. Thus the outputted sub-CAD is not a variety sub-CAD but a superset of cells containing one. Algorithm 2 demonstrates this approach.
(b) Use Collins-Hong projection [28] for the first $(n-k)$ projection stages. Then apply Algorithm 1 with McCallum's reduced projection operator for equational constraints to build a variety sub-CAD of $\mathbb{R}^k$ (as verified by Theorem 1) before continuing lifting to a variety sub-CAD of $\mathbb{R}^n$.

The latter approach is still a variety sub-CAD and allows for a smaller CAD of $\mathbb{R}^k$ but these benefits may be overshadowed in the final sub-CAD of $\mathbb{R}^n$ due to lifting with respect to a larger set of polynomial in the later stages. Of course both cases may return FAIL in which case having all sub-algorithms implement [28] would be the only way forward. However, it is worth noting that when the equational constraint is not in the main variable we may actually avoid unnecessary failure: if the input is not well-oriented but the problematic nullification only occurs on cells that are not on the variety then the outputted sub-CAD will still be valid. It is interesting to note that this sub-CAD is now a subset of a CAD we do not know how to produce algorithmically.

Algorithm 2: Algorithm to produce sub-CADs with respect to a variety of lower dimension.

Input : A formula φ, a declared equational constraint $f = 0$ from φ and variables $\boldsymbol{x} = x_1, \ldots, x_n$. φ is in $\boldsymbol{x}$ and all factors of f have the same main variable.

Output: A sub-CAD $\mathcal{D}$ on which φ is truth invariant and which is the superset of a variety sub-CAD for (φ, f), or FAIL.

```
1  Extract from φ the set of polynomials A;
2  Set k to be the index of mvar(f) ;
3  Perform the first n − k projection stages using ProjOp and starting with A.
4  Set P_1 to be the projection polynomials with main variable x_i
5  if k ≠ 1 then
6      D_{k−1} ← CADAlgo(P_{k−1}, [x_1, ..., x_{k−1}]) ;        // Computation of a CAD of ℝ^{k−1}
7  if D_{k−1} = FAIL then
8      return FAIL ;                                            // P_{k−1} is not well oriented
9  D_k ← [] ;
10 if k = 1 then
11     Set S to be the CAD formed by decomposing ℝ according to the roots of P_1.;
12     if |S| > 1 then
13         for i = 1 ... (|S| − 1)/2 do
14             D_k.append(S[2i]) ;                              // Cells with even index included
15 else
16     for c ∈ D_{k−1} do
17         S ← GenerateStack(P_k, c);                           // kth lifting stage
18         if S = FAIL then
19             return FAIL ;                                    // P_k is not well oriented
20         if |S| > 1 then
21             for i = 1 ... (|S| − 1)/2 do
22                 D_k.append(S[2i]) ;                          // Cells with even index included
23 for i = k + 1, ... n do
24     D_i ← [];
25     for c ∈ D_{i−1} do
26         S ← GenerateStack(P_i, c);                           // ith lifting stage
27         if S = FAIL then
28             return FAIL ;                                    // P_i is not well oriented
29         D_i.append(S) ;                                      // All cells included
30 return D_n;
```

Example 2 demonstrates the savings offered by a variety of lower dimension, but also the difficulty in obtaining an actual V-sub-CAD.

Example 2.
Consider again the polynomials from Example 1 (with the same variable ordering) but this time with the formula $\varphi_2 := f < 0 \wedge g = 0$. There is only one cell where this is true (the y-axis inside the circle). A minimal V-sub-CAD truth-invariant for φ_2 would contain only the 5 cells shown on the left of Figure 2. However, if we use Algorithm 2 then we would produce 11 cells, as shown on the right. Algorithm 2 first builds a CAD of $\mathbb{R}^1$ with respect to all the univariate projection polynomials. This has 7 cells (the points $-1, 0, 1$ and the intervals in-between). It discards the intervals and lifts over the

three points with respect to f obtaining the 11 cells shown. No lifting (and hence real root isolation) was performed over the other 4 cells in $\mathbb{R}^1$ since we could conclude they were not part of the variety.

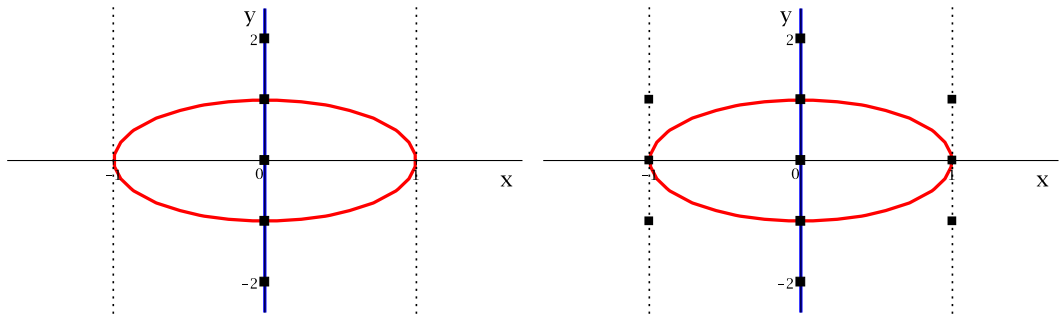

FIGURE 2. Images representing the CADs described in Example 2. The minimal variety sub-CAD is on the left with the output from Algorithm 2 on the right.

To allow factors of the equational constraint with different main variable would require a further extension to perform multiple stages of restricted lifting (steps $9-22$ of Algorithm 2) when lifting to $\mathbb{R}^i$ where x_i is a main variable of a factor, and full lifting (steps $23-29$) otherwise.

Finally, note that we have only discussed using a single (designated) equational constraint. In the case of two or more (as in $f_1 = 0 \wedge f_2 = 0 \wedge \ldots$) we need the theory of bi-equational constraints and beyond. Although there has been some work on this [35, 13] we have not investigated its interaction with variety sub-CADs yet.

2.2. Layered sub-CADs

The idea of returning CAD cells of certain dimensions was first discussed in [31] and revisited in [40, 10]. All these papers discuss the idea of returning only CAD cells of full-dimension, noting that this is sufficient to solve problems involving only strict polynomial inequalities. This idea was extended in the technical report [45] to the case of returning cells with dimension above a prescribed value. We reproduce some of that unpublished work here, including the key algorithms.

Definition 4.
Define the set of cells in a CAD of a given dimension as a **layer**. Let ℓ be an integer with $1 \leq \ell \leq n+1$. Then an **ℓ-layered Sub-CAD (ℓ-L-sub-CAD)** is the subset of a CAD of dimension n consisting of all cells of dimension $n-i$ for $0 \leq i < \ell$. We refer to a CAD consisting of all cells of all dimensions as a **complete CAD**.

Remark 2.

1. An ℓ-layered CAD consists of the top ℓ layers of cells in a CAD. The dimensions of these cells will depend on the dimension of the space the CAD decomposes.
2. In the literature the set of cells of full-dimensional has been referred to as an open CAD, a full CAD and a generic CAD. We prefer layered CAD as it is less open to misinterpretation and allows us to generalise the idea beyond the top dimension. The set of cells of full dimension is then a 1-layered CAD. Note also that a complete CAD of $\mathbb{R}^n$ has $(n+1)$-layers.
3. When building a 1-layered CAD it was pointed out in [40] that a simplified projection operator could be used. Instead of taking the full set of coefficients for a polynomial only the leading coefficient is required (since the others are there to ensure delineability if the first vanishes, but this could only happen on a cell of less than full dimension). In this paper we focus on improvements to the lifting phase, but if only a 1-layered CAD is required then this further saving in the projection phase is available.

Algorithm 3 describes how an ℓ-layered CAD may be produced. The main idea is that during the lifting process we check cell dimension before each stack construction. If a cell has too low a dimension to give cells in $\mathbb{R}^n$ of dimension ℓ or higher then it is discarded. We give a general algorithm which can be used with any suitable (and compatible) projection operator and stack generation procedure.

Algorithm 3: `LayeredSubCAD`$(\varphi, \ell, \mathbf{x})$: Algorithm to produce ℓ-layered sub-CADs.

Input : A formula φ, an integer $1 \leq \ell \leq n+1$ and variables $\boldsymbol{x} = x_1, \dots x_n$. φ is in $\boldsymbol{x}$.
Output: An ℓ-layered sub-CAD for φ, or FAIL.

1 $\mathbf{P} \leftarrow$ output from applying `ProjOp` repeatedly to φ ; `// Full projection phase`
2 **for** $i = 1, \dots, n$ **do**
3 Set $\mathbf{P}[i]$ to be the projection polynomials with mvar x_i;
4 Set $\mathcal{D}[1]$ to be the CAD of $\mathbb{R}^1$ obtained by isolating the roots of $\mathbf{P}[1]$;
5 **for** $i = 2, \dots, n$ **do**
6 $\mathcal{D}[i] \leftarrow [\,]$;
7 **for** $c \in \mathcal{D}[i-1]$ **do**
8 `dim` $\leftarrow \sum_{\alpha \in c.\texttt{index}} (\alpha \mod 2)$; `// Lift over suitable dimension cells`
9 **if** `dim` $> i - \ell - 1$ **then**
10 $S \leftarrow$ `GenerateStack`$(\mathbf{P}[i], c)$;
11 **if** $S =$FAIL **then**
12 **return** FAIL ; `// Input is not well oriented`
13 **else**
14 $\mathcal{D}[i]$.`append`(S)
15 $\mathcal{D} \leftarrow [\,]$;
16 **for** $c \in \mathcal{D}[n]$ **do**
17 `dim` $\leftarrow \sum_{\alpha \in c.\texttt{index}} (\alpha \mod 2)$;
18 **if** `dim` $> n - \ell$ **then**
19 $\mathcal{D}$.`append`(c); `// Remove cells of low dimension from final lift`
20 **return** $\mathcal{D}$;

Theorem 2. *Algorithm 3 satisfies its specification.*

Proof. The output being a subset of a valid CAD follows from the correctness and compatibility of the sub-algorithms used (as proven in [17, 28, 33, 34, etc.]). It remains to verify that the cells discarded could not contribute cells in the top ℓ layers of the outputted sub-CAD.

When lifting over a cell of dimension d in a sub-CAD of $\mathbb{R}^i$ it can contribute cells in the sub-CAD of $\mathbb{R}^n$ of dimension at most $d + n - i$. If these are required for an ℓ-layered sub-CAD they must have dimension at least $n - \ell$, and so we can discard them if $d \leq i - \ell - 1$ (step 9).

When performing the final lift we must build stacks over cells of dimension $n - \ell - 1$ or greater, but of course some of the cells in those stacks may not have dimension $n - \ell$. Hence we check for this at the end (step 18) only keeping those of the required dimension. □

Example 3.
Consider once again the polynomials introduced by Example 1 (with variable ordering $x \prec y$). The first image in Figure 1 demonstrated that the simplest sign-invariant complete CAD for the polynomials would have 23 cells. Figure 3 shows the same CAD, this time with the dimensions of each cell indicated.

Suppose that the polynomials arise from a problem involving the formula $\varphi_3 := f < 0 \wedge g < 0$. The solution is then given by the single cell inside the circle to the left of the y-axis. In this case we know the solutions must all be cells of full-dimension and thus that a 1-layered CAD would suffice. Algorithm 3 would return only the 8 cells of dimension 2 (the boxes in Figure 3).

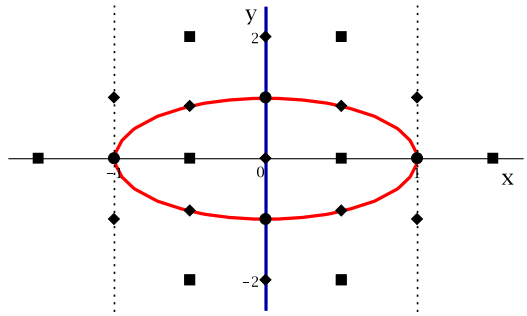

FIGURE 3. Image representing the CAD described in Example 3. The boxes indicate cells of dimension 2, the diamonds dimension 1 and the spheres dimension 0.

In Example 3 a 1-layered sub-CAD is φ-sufficient. This was a case where the problem consisted of strict inequalities meaning the solutions have full-dimension (as noted previously in [31, 40]). More generally there are classes of problems with known solution dimension for which the layered sub-CAD technology will be beneficial. For example, recall the cyclic n-polynomials (the set of n symmetric polynomials in n variables). In [2] it is shown that if there exists $m > 0$ such that $m^2|n$ then there are an infinite number of roots and they are, at least, of dimension $(m-1)$. A sub-CAD containing cells of dimension $(m-1)$ and higher would therefore be sufficient to identify the largest families of solutions, and so a layered sub-CAD could be an appropriate tool.

Layered sub-CADs will also be useful when, although a sub-CAD may not be φ-sufficient, the underlying application only requires generic families of solutions. For example, when robot motion planning we need only identify paths through cells of full-dimension, and so a 1-layered sub-CAD may be enough. We consider a motion planning example later in Section 5.3. In that example we actually need the full-dimensional cells on a variety, which are part of the 2-layered sub-CAD. We find that the 1-layered variety sub-CADs developed in Section 3.1 are also appropriate.

We now explain how Algorithm 3 may be adapted to a recursive procedure. Our general principle when constructing layered sub-CADs is to stop lifting over a cell if it cannot lead to cells of sufficient dimension in $\mathbb{R}^n$. This occurs when the cell in question was produced as a section, rather than a sector. Let us call these **terminating sections**. Instead of discarding these sections after they are produced, the recursive algorithm stores them in a separate output variable for use later if required.

For example, when constructing a 1-layered CAD these terminating sections will have dimension $m-1$ where m is the number of variables lifted at the level the section is constructed (i.e. they are deficient by one dimension). So at the final lifting stage the terminating sections will have dimension $n-1$. Now suppose we wish to extend to a 2-layered CAD. If we take one of these terminating sections of any dimension and construct successive stacks over it we will obtain cells that are of dimension $n-1$ and terminating sections that are deficient (with respect to their level) by two dimensions. Combining these $n-1$-dimensional cells with the 1-layered CAD produces a 2-layered CAD. If we again store the remaining (and new) terminating sections then they can be used later to construct a 3-layered CAD if desired. In this manner we can recursively produce ℓ-layered CADs, useful if the numbers of layers required is not known at the start of the computation. The method is described by Algorithm 4.

Theorem 3. *Algorithm 4 satisfies its specification.*

Proof. As with the proof of Theorem 3 the correctness mostly follows from the correctness of the sub-algorithms used and it remains only to verify that the correct layers are produced. We see that if $\mathcal{LD}$ is empty then a CAD of the real line is produced and the sections and sectors are separated. Otherwise, each of the cells in the existing terminating sections are lifted over. The stacks constructed consist of sectors, which together with $\mathcal{LD}$ form a layered sub-CAD with one extra layer, or new terminating sections which could be lifted over to gain the complete CAD. □

Algorithm 4: `LayeredSubCADRecursive`$(\varphi, \mathbf{x}, \mathcal{C}, \mathcal{LD})$: Algorithm that may be applied recursively to produce layered sub-CADs with increasing numbers of layers.

Input : A formula φ, variables $\boldsymbol{x} = x_1, \ldots, x_n$, a list of cells of $\mathbb{R}^n$ labelled $\mathcal{LD}$, and a list of lists of cells, $\mathcal{C}$ (ordered by increasing dimension). The formula φ is in $\boldsymbol{x}$. The lists $\mathcal{C}$ and $\mathcal{LD}$ may be empty. Otherwise the list $\mathcal{LD}$ must contain a layered sub-CAD for φ and $\mathcal{C}$ the corresponding terminating sections.

Output: Either FAIL or a layered sub-CAD $\mathcal{D}'$ for φ. If $\mathcal{LD}$ was empty then $\mathcal{D}$ is a 1-layered sub-CAD and otherwise it is a layered sub-CAD with including cells of dimension one layer lower than $\mathcal{LD}$. Also, an updated list of lists of cells containing the terminating sections sufficient to construct the complete CAD.

```
global P ;                                  // To avoid recomputing projection polynomials
if P is undefined then
    P ← output from applying ProjOp repeatedly to φ ;        // Full projection phase
    for i = 1, ..., n do
        Set P[i] to be the projection polynomials with mvar x_i;
C' ← [ ];
if C == ∅ then
    D ← [ ] ;                                                // Base case - construct ℝ^1
    Set Base to be the CAD of ℝ^1 obtained by isolating the roots of P[1];
    for i = 1, ..., length(Base) do
        if (i mod 2) == 1 then
            D[1].append(Base[i]) ; // Sectors only (odd index) for the 1-layered CAD
        else
            C'[1].append(Base[i]) ;              // Sections (even index) added stored
    D' ← [ ];
else
    D ← C ;                                 // Use previously computed terminating sections
    D' ← C[n];
for i = 2, ..., n do
    for c ∈ D[i − 1] do
        S ← GenerateStack(P[i], c);
        if S = FAIL then
            return FAIL ;                                    // Input is not well oriented
        for j = 1, ..., length(S) do
            if (j mod 2) == 1 then
                D[i].append(S[j]) ;          // Sector (odd index) so add to output CAD
            else
                C'[i].append(S[j]) ;                     // Section (even index) so store
D' ← D' ∪ LD ;                  // Combine new cells with those previously computed
return [D', C'];
```

Algorithm 4 has been implemented in Maple. It uses a global variable to avoid recalculation of projection polynomials, and Maple's unevaluated function call syntax (prefixing the command with %) allowing it to be repeatedly evaluated to return sub-CADs with an increasing number of layers.

We finish this subsection by noting the following interesting property of certain layered CADs.

Theorem 4 ([45]).
Let $F \subset \mathbb{Z}[x_1, \ldots, x_n]$ and $\mathcal{D}$ be a 1- or 2-layered CAD of $\mathbb{R}^n$ sign-invariant for F. Then $\mathcal{D}$ is order-invariant with respect to F, meaning each polynomial has constant order of vanishing on each cell.

Order-invariance is a stronger property than sign-invariance, but the extra knowledge it gives allows for the validated use of smaller projection operators (see for example [33]). Hence this property allows for the avoidance of well-orientedness checks during stack generation when building 1 or 2-layered sub-CADs. This means not just a savings in computation time but the avoidance of unnecessary failure declarations that can sometimes follow from such checks (as discussed in [9, 24]).

2.3. Other sub-CADs in the literature

We note that other ideas from the CAD literature could be described as, or be easily adapted to produce, sub-CADs. We list these for completeness.

In [32], whilst trying to solve a motion planning problem in the plane described by a formula φ, the author identifies a subset of cells in the decomposition of $\mathbb{R}^1$ for which any valid cell for φ must lie over. Lifting over only these cells only gives a φ-sufficient sub-CAD. Similar ideas are in [29].

In [10] an algorithm is presented which given polynomials F and a point α, returns a single cell containing α on which F is sign-invariant. The cell belongs to a CAD (although not necessarily one that could be produced by any known algorithm) and hence this is an extreme example of a sub-CAD.

Partial CAD [18] works by avoiding the splitting of cylinders into stacks when lifting over cells where the truth value is already known. If cells over which the truth value is false were instead simply discarded then what is left would be a sub-CAD sufficient to analyse the input formula.

In [39] an algorithm is described which takes polynomials F and returns a CAD D and theory Θ (set of negated equations). The CAD is sign-invariant for F for all points which satisfy Θ. Rather than a sub-CAD of $\mathbb{R}^n$ this is actually a CAD of $\mathbb{R}^n_\Theta$: all those points in $\mathbb{R}^n$ except the set of measure zero which do not satisfy Θ.

In [43], an algorithm is given for solving systems over cylindrical cells described by cylindrical algebraic formulae. This allows cells produced from a CAD to be used easily in further computation, which may implicitly be used to produce either sub-CADs or CADs of a sub-space.

3. Extending the use of sub-CAD

Layered and variety sub-CADs can offer significant savings individually but this can be increased by combining them with each other, as discussed in Section 3.1. We then consider how they may interact with the recent theory of truth-table invariant CAD in Section 3.2.

3.1. Combining layered and variety sub-CADs

The idea behind sub-CADs is simple: we wish to filter out only those cells of relevance to us. A variety sub-CAD does this in one step during a single stage of the lifting phase, whereas a layered sub-CAD stratifies the cells throughout the whole lifting process. There is no reason why these two ideas cannot be combined. We discuss this in the case where all factors of the equational constraint defining the variety have main variable x_n and are not nullified on a low-dimensional cell. (It may be generalised but this would require caution as discussed in Section 2.1.) The key is the following simple result.

Lemma 1.
Let $\mathcal{D}$ be a CAD of $\mathbb{R}^{n-1}$ for some formula φ and let c be a cell in $\mathcal{D}$ of dimension k. Further, suppose $f = 0$ is an equational constraint of φ, each factor of f has main variable x_n, and f is not nullified on c. Then any section of the stack lifted over c with respect to f will have dimension k.

The lemma shows that lifting over a cell onto a variety gives cells with the same dimension. Hence when lifting over an ℓ-layered sub-CAD of $\mathbb{R}^{n-1}$ (which contains cells of dimension $n-\ell, \ldots, n-1$)

we produce a sub-CAD containing all the cells of dimensions $n-\ell,\dots,n-1$ on the variety. We may think of this as an ℓ-layered sub-CAD *of the variety*.

Definition 5.
Let φ be a Tarski formula with equational constraint $f=0$ which has main variable x_n in all its factors, and let $1\leq\ell\leq n$. A truth-invariant sub-CAD for φ whose cells have dimension $n-i-1$ for $0\leq i<\ell$ and rest on the variety defined by $f=0$ is an **$\boldsymbol{\ell}$-Layered Variety Sub-CAD ($\boldsymbol{\ell}$-LV-sub-CAD)**.

Remark 3.
Note that in general an ℓ-layered variety sub-CAD consists of the top ℓ layers of cells *on the variety*. This can be thought of as the intersection of an $(\ell+1)$-layered CAD of $\mathbb{R}^n$ with the variety (as the layer of n-dimensional cells is discarded when lifting to the variety).

Lemma 1 leads to Algorithm 5 for producing $\boldsymbol{\ell}$-LV-sub-CADs.

Algorithm 5: `LayeredVarietySubCAD`$(\varphi, f, \ell, \boldsymbol{x})$: Algorithm for ℓ-layered variety sub-CADs.

Input : A formula φ, a declared equational constraint $f=0$ from φ, an integer $1\leq\ell\leq n+1$ and variables $\boldsymbol{x}=x_1,\dots,x_n$. φ is in $\boldsymbol{x}$ and all factors of f have main variable x_n.
Output: An ℓ-layered variety sub-CAD $\mathcal{D}$ for φ, or FAIL.

```
1  Extract from φ the set of polynomials A and from f the subset E ⊂ A ;
2  P ← output from applying ProjOp to (A,E) ;                  // First projection stage
3  D' ← LayeredSubCAD(P, ℓ) ;                                  // Computation of a sub-CAD of R^{n-1}
4  if D' = FAIL then
5      return FAIL ;                                           // P is not well oriented
6  D ← [];
7  for c ∈ D' do
8      S ← GenerateStack(E, c);                                // Final lifting stage
9      if S = FAIL then
10         return FAIL ;                                       // Input is not well oriented
11     if |S| > 1 then
12         for i = 1 ... (|S| − 1)/2 do
13             D.append(S[2i]) ;                               // Cells with even index are sections
14 return D;
```

Theorem 5. *When the sub-algorithms are chosen to implement McCallum's algorithm to produce CADs with respect to an equational constraint* [34]*, then Algorithm 5 satisfies its specification with the outputted sub-CAD consisting of cells on which the input formula has constant truth value.*

Proof. As with Theorem 1 the sub-CAD structure and invariance property follow from the results in [34]. Theorem 1 verifies that $\mathcal{D}$ is an ℓ-layered sub-CAD of $\mathbb{R}^{n-1}$ and Lemma 1 concludes that the lifting in the final loop results in an ℓ-layered sub-CAD of $\mathbb{R}^n$. The case where Lemma 1 does not hold is a case where the input is not-well oriented and thus FAIL is returned. □

We now give a simple example which uses both the layered and variety sub-CAD ideas together, as well as illustrating the difference variable ordering can make.

Example 4.
Consider a final time the polynomials from Example 1 and this time the formula $\varphi_1 := f=0\wedge g<0$. In Example 1 we saw that with variable ordering $x\prec y$ a variety sub-CAD could be produced with 8 cells, 3 of which described the solutions. Instead let us build a 1-layered variety sub-CAD as represented in the first image of Figure 4. The output would now be only 4 cells (those indicated with solid boxes).

Two of these describe the generic solution sets $\{x \in (-1, 0), y = \pm\sqrt{1 - x^2}\}$ but the third and final cell in the solution set $\{x = -1, y = 0\}$ has been lost. Note that in this case (unlike the variety sub-CAD) there will be a reduction in CAD computation time as there will be no lifting over cells of dimension zero in the CAD of $\mathbb{R}^1$. The cells which have been computed but discarded are shown by empty boxes.

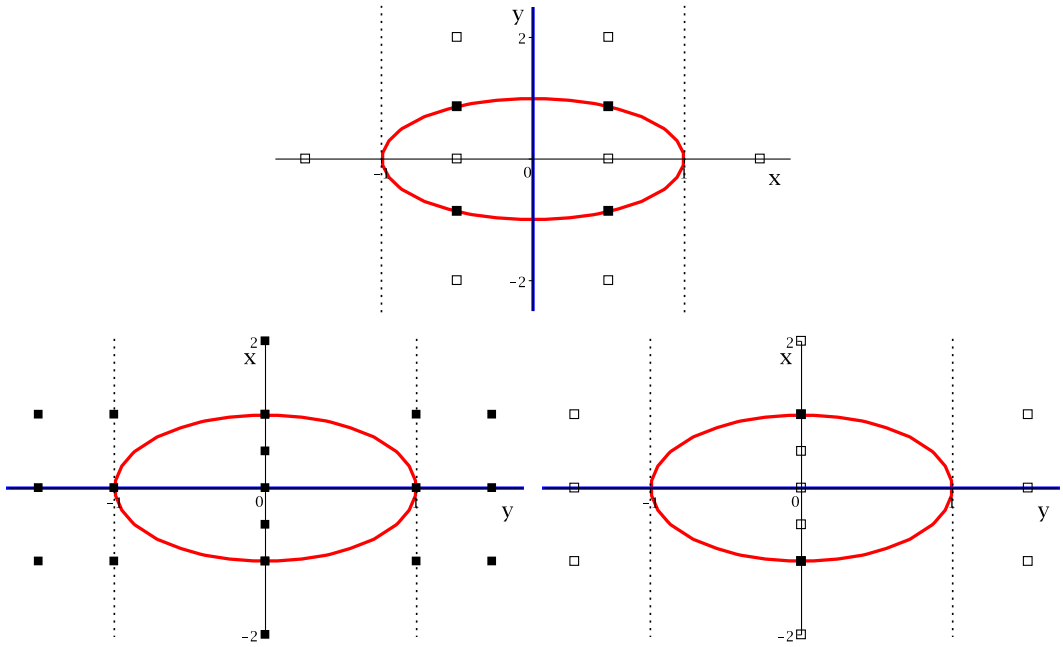

FIGURE 4. Figure representing the CADs described in Example 4. The solid boxes represent cells returned and the empty boxes cells computed but discarded. The first image uses variable ordering $x \prec y$ and the others $y \prec x$ (axes labelled accordingly).

Consider now the alternative variable ordering $y \prec x$. In this case a minimal complete sign-invariant CAD must have 19 cells, as represented by the second image in Figure 4. A variety sub-CAD would return the 4 cells describing the circle and a 1-layered variety sub-CAD only the 2 cells of these which have full-dimension on the variety. This sub-CAD is represented by the final image in Figure 4. Again, there has been a saving in CAD construction time but in this case the cells returned include the full solution set (the single cell describing the upper half of the circle).

3.2. Truth table invariant sub-CADs

The new ideas of variety and layered sub-CADs can be further adapted by combining with other CAD techniques. We shall discuss here the interaction with the idea of truth table invariant CAD (TTICAD) introduced recently by [4]. We recall the definition of a TTICAD.

Definition 6 ([4])**.**
Let $\Phi := \{\phi_i\}_{i=1}^t$ be a list of quantifier-free formulae (QFFs). We say a cylindrical algebraic decomposition $\mathcal{D}$ is a **Truth Table Invariant** CAD for Φ (a **TTICAD**) if the Boolean value of each ϕ_i is constant (either true or false) on each cell of $\mathcal{D}$.

In [4] an algorithm was given to build TTICADs in the case where each ϕ_i contained an explicit equational constraint. This has recently been extended to allow for any ϕ_i [25]. The algorithms involve a new projection operator which encapsulates the interaction between the equational constraints of the ϕ_i's whilst ignoring interactions between polynomials that have no effect on the truth value of ϕ_i.

Truth table invariance is very useful. Given a problem defined by a parent formula Φ built by a boolean combination of $\{\phi_i\}_{i=1}^t$, a TTICAD is both sufficient to determine where Φ is true, and

more efficient than any other projection operator. Further, there are classes of problems for which a TTICAD is exactly the desired structure, such as the problem of decomposing a complex domain according to the branch cuts of multivariate functions [3, 37, 26].

In the case where there is a parent formula Φ and all the ϕ_i have their own equational constraint there exists a variety on which the solution rests. It is defined by the product of the individual equational constraints (an implicit equational constraint for Φ). Hence in this case it makes sense to build a **Variety Sub-TTICAD (V-sub-TTICAD)**. For simplicity, we assume that each equational constraint has factors which all have main variable x_n. We can then use Algorithm 1 on Φ with the declared equational constraint the product of the individual ones. `ProjOp` should be the TTICAD projection operator from [4] applied to the sequence of bases of polynomials appearing in the formulae and `GenerateStack` an algorithm which checks for the TTICAD well-orientedness properties. A proof of the validity of Algorithm 1 using these sub-procedures to give a V-sub-TTICAD would follow analogously to the proof of Theorem 1. Although the variety considered is an implicit equational constraint the TTICAD projection theory is more efficient that the equational constraint theory applied to this case (as described in [5]).

Creating a **Layered Sub-TTICAD (L-sub-TTICAD)** is also possible, by applying Algorithm 3 (or 4) with the TTICAD projection operator and a `GenerateStack` algorithm which checks for its well-orientedness condition. The validity is proven by Theorem 2 (using results from [4]).

Similarly, we can combine TTICAD with the ideas of Section 3.1 to produce layered variety sub-TTICADs as defined below.

Definition 7.
Let $\{\phi_i\}_{i=1}^t$ be a list of QFFs with each ϕ_i having equational constraint f_i whose factors all have main variable x_n. Let $1 \leq \ell \leq n$. A sub-CAD of a TTICAD for $\{\phi_i\}_{i=1}^t$ containing all cells of dimension $n-1-i$ for $0 \leq i < \ell$ resting on the variety defined by $\prod_{i=1}^t f_i = 0$ is called a **Layered Variety Sub-TTICAD (LV-sub-TTICAD)**.

Remark 4.
As with ML-sub-CADs, a ℓ-layered variety sub-TTICAD consists of the top ℓ-layers of cells *on the variety* $\prod_{i=1}^t f_i = 0$. Again, this can be thought of as the intersection of an $(\ell+1)$-layered CAD of $\mathbb{R}^n$ with the variety (as the layer of n-dimensional cells is discarded when lifting to the variety).

We can construct LV-sub-TTICADs by using Algorithm 5 with the sub-algorithms implementing TTICAD. As noted, the correctness of the approaches from this subsection follow analogously to the proofs of Theorem 1 - 3 and 5 (noting that the exceptional cases where part of a variety is nullified would have meant the input was not well-oriented for TTICAD and thus triggering an output of FAIL). The example in Section 5.2 demonstrates the use of a LV-sub-TTICAD and the benefits of choosing to do so. See the technical report [45] for some further details and examples.

4. Complexity analysis

We provide a complexity analysis of the algorithms to compute sign-invariant variety sub-CADs and 1-layered variety sub-CADs in the case where the equational constraint has all factors with main variable x_n. We need to study three parts of the complexity:

- **Projection.** The complexity of the equational constraint projection set needs to be analysed. In particular the number of polynomials, the maximum degree and size of their coefficients.
- **Calculation of $(n-1)$ dimension CAD.** These values then can be used to estimate the complexity of the $(n-1)$–dimensional CAD.
- **Lifting.** Finally the complexity of the lifting stage can be combined with the previous step to describe the complexity of the variety CAD.

We recall the previous comprehensive work on CAD complexity by Collins [17] and McCallum [31].

We first standardise some notation for a CAD with respect to a set of polynomials A: let n be the number of variables, m the number of polynomials in A, d the maximum degree in any variable

of the polynomials in A, and l the maximum norm length of the polynomials in A (where the norm length, $|f|_1$, is the sum of the absolute values of the integer coefficients of a polynomial).

Let $A_1 := A$ and let $A_{i+1} = \text{Proj}(A_i)$. In general the projection operator used will be clear: most of the following is with respect to Collins' projection operator and, therefore, is a 'worst case scenario' compared to the improved operators of McCallum [33] and Brown [7] (which are subsets of the Collins operator). Let m_k be the number of polynomials in A_k, d_k the maximum degree of A_k, and l_k the maximum norm length.

4.1. Collins' algorithm

Collins [17] works through the original CAD algorithm in great detail to analyse the complexity, and this methodology is followed in [31]. We recall some key results, noting they could all be uniformly improved with ideas from papers such as [14, 19] (but that is not the aim of this paper).

In the projection stage we can bound the properties of the projection sets as follows:

$$m_k \leq (2d)^{3^k} m^{2^{k-1}}; \quad d_k \leq \frac{1}{2}(2d)^{2^{k-1}}; \quad l_k \leq (2d)^{2^k} l.$$

By combining these bounds Collins shows the projection phase is dominated by

$$(2d)^{3^{n+1}} m^{2^n} l^2.$$

The base case and lifting algorithm requires the isolation of real roots of univariate polynomials. Collins bounds this procedure as follows. Let A be a set of univariate polynomials with degree bounded by d and norm length bounded by l. Then for a given $f \in A$ with $\hat{d} := \deg(f)$ and $\hat{l} := |f|_1$ a lower bound on the distance between two roots is given by

$$\frac{1}{2}\left(\sqrt{e}\hat{d}^{\frac{3}{2}}\hat{l}\right)^{-\hat{d}}. \tag{1}$$

Collins uses his analysis of Heindel's algorithm for real root isolation to show that isolating the roots is dominated by

$$\hat{d}^8 + \hat{d}^7\hat{l}^3. \tag{2}$$

Therefore the operations needed to isolate all roots in A, with $m := |A|$, is dominated by:

$$md^8 + md^7l^3. \tag{3}$$

For a given h, refining a root interval to 2^{-h} is dominated by $d^2h^3 + d^2l^2h$. Collins multiplies all polynomials in A and uses (1) to show that all roots are separated by:

$$\delta := \frac{1}{2}\left(\sqrt{e}(md)^{\frac{3}{2}}l^m\right)^{-md}.$$

If we take h to be $\log(\delta)$ we can refine all intervals of the polynomials in

$$md(d^2(m^2dl + md\log(md))^3 + md^3l^3) = O(m^7d^7l^3 + m^2d^4l^3).$$

Combining this with (3) tells us that the necessary refinement of intervals for all the polynomials is dominated by

$$md^8 + m^7d^7l^3. \tag{4}$$

Combining (4) with the size of the full projection set concludes the base phase is dominated by

$$(2d)^{3^{n+3}} m^{2^{n+2}} l^3.$$

We also need to consider the polynomials involved in the lifting stage. This involves looking at the univariate polynomials created after substituting sample points, along with the polynomials required to define the algebraic extensions of $\mathbb{Q}$ that those sample points are contained in.

We follow [17, 31] in using primitive elements to calculate the costs of operations, even though implementers are extremely unlikely to use them. For each sample point $\beta = (\beta_1, \ldots, \beta_k) \in \mathbb{R}^k$ there is a real algebraic number $\alpha \in \mathbb{R}$ such that $\mathbb{Q}(\beta_1, \ldots, \beta_k) = \mathbb{Q}(\alpha)$. Let A_α be the polynomial in $\mathbb{Q}[x]$ that, along with an isolating interval I_α, defines α (so $A(\alpha) = 0$). Let d_k^* be the maximum degree of

these A_α, and l_k^* the maximum norm length. Each coordinate β_i is represented in $\mathbb{Q}(\alpha)$ by another polynomial. Let l'_k be the maximum norm length of these polynomials. Then

$$d_k^* \leq (2d)^{2^{2n-1}}, \quad \text{and} \quad l_k^*, l'_k \leq (2d^2)^{2^{2n+3}} m^{2^{n+1}} l.$$

Let u_k be the number of univariate polynomials (after substitution) and c_k the number of cells at level k. Then

$$u_k, c_k \leq 2^{2^n} \prod_{i=1}^{n} m_i d_i, \quad \text{and} \quad u_k, c_k \leq (2d)^{3^{n+1}} m^{2^n}. \tag{5}$$

Collins combines all these results to give a complexity bound for the full algorithm of

$$(2d)^{4^{n+4}} m^{2^{n+6}} l^3. \tag{6}$$

4.2. McCallum's CADMD algorithm

In [31], McCallum gives a complexity bound, using the same methodology as [17], for his CADMD algorithm which produces a 1-layered CAD. Thanks to the avoidance of algebraic numbers (all sample points in a 1-layered CAD can be produced directly in $\mathbb{Q}$) the exponents are lower than in (6). The complexity is dominated by

$$(2d)^{3^{n+4}} m^{2^{n+4}} l^3. \tag{7}$$

4.3. Analysis of $P_E(A)$ and $(n-1)$-dimensional CADs

To have an accurate complexity we must consider some properties of the projection set: the size, maximum degree, and maximum norm length. Let E be the subset of A containing the factors of the designated equational constraint. Let $m_A := |A|$, $m_E := |E|$, $m_{A\setminus E} := |A \setminus E|$, and let d_A, d_E, $d_{A\setminus E}$, l_A, l_E, and $l_{A\setminus E}$ be defined similarly.

Since we are constructing a CAD with respect to equational constraints, we assume use of the projection operator, $P_E(A)$, defined by McCallum [34] as

$$P_E(A) := P(E) \cup \{\text{res}_{x_n}(f, g) \mid f \in E, g \in A \setminus E\}$$

where $P(E)$ is an application of the operator defined in [33] (giving the coefficients, discriminants and cross resultants of E). We will denote the resultant set, $P_E(A) \setminus P(E)$, by $\text{ResSet}_\text{E}(\text{A})$.

The size of $P(E)$ is bounded by the number of coefficients ($m_E d_E$), discriminants (m_E), and resultants ($\binom{m_E}{2} = \frac{m_E(m_E-1)}{2}$). Therefore we have

$$\begin{aligned} |P_E(A)| \leq m_E d_E + m_E + \frac{m_E(m_E-1)}{2} + m_E m_{A\setminus E} &= \frac{m_E}{2}\left(2d_E + 2m_{A\setminus E} + m_E - 1\right) \\ &= \frac{m_E}{2}\left(2d_E + m_{A\setminus E} + m_A - 1\right). \end{aligned} \tag{8}$$

The maximum degree of $P_E(A)$ is the greater of the maximum degrees of $P(E)$ and the resultant set. We also have a bound on the degree of a resultant with respect to x:

$$\deg(\text{res}_x(f, g)) \leq (\deg_x f + \deg_x g) \cdot (\max(\deg_y f, \deg_y g)).$$

Using our overall degree bounds gives

$$\begin{aligned} \max\deg(\text{ResSet}_E(A)) \leq (d_E + d_{A\setminus E}) \cdot \max(d_E, d_{A\setminus E}) &= \max(d_E^2 + d_E d_{A\setminus E}, d_{A\setminus E}^2 + d_E d_{A\setminus E}) \\ &\leq \max(2d_E^2, 2d_{A\setminus E}^2). \end{aligned} \tag{9}$$

We also have

$$\max\deg(P(E)) \leq \max(d_E, 2d_E^2, 2d_E^2) = 2d_E^2. \tag{10}$$

Combining (9) and (10) gives a degree bound for $P_E(A)$:

$$\max\deg(P_E(A)) \leq \max(2d_E^2, 2d_{A\setminus E}^2) \leq d_A^2. \tag{11}$$

Finally, if we denote the maximum norm length of $P_E(A)$ by $\overline{l}$, then we know $\overline{l} \leq l_2$ (since $P_E(A) \leq \text{Proj}(A)$) and so

$$\overline{l} \leq l_2 \leq (2d_A)^{2^2} l_A = 16 d_A^4 l_A. \tag{12}$$

Substituting the bounds from (8), (11) and (12) into (6) and (7) gives an estimate on the complexity of a complete $(n-1)$-dimensional $P_E(A)$-invariant Collins CAD dominated by

$$\leq \left(2 \cdot 2d_A^2\right)^{2^{2(n-1)+8}} \left(\frac{m_E(2d_E + m_A + m_{A\setminus E} - 1)}{2}\right)^{2^{n-1+6}} \left(16 d_A^4 l_A\right)^3$$

$$\leq 16^3 \left(4d_A^2\right)^{2^{2n+6}} \left(\frac{m_E(2d_E + 2m_A - 1)}{2}\right)^{2^{n+5}} l_A^3 d_A^{12}. \tag{13}$$

Similarly, the complexity of a 1-layered $(n-1)$-dimensional $P_E(A)$-invariant CAD is dominated by

$$\leq \left(2 \cdot 2d_A^2\right)^{3^{n-1+4}} \left(\frac{m_E(2d_E + m_A + m_{A\setminus E} - 1)}{2}\right)^{2^{n-1+4}} \left(16 d_A^4 l_A\right)^3$$

$$\leq 16^3 \left(4d_A^2\right)^{3^{n+3}} \left(\frac{m_E(2d_E + 2m_A - 1)}{2}\right)^{2^{n+3}} l_A^3 d_A^{12}. \tag{14}$$

4.4. Overall complexities

So far, we have computed the complexities of the $(n-1)$-dimensional CADs. The final step is to lift over these cells with respect to the equational constraints. From (5) we can bound the number of univariate polynomials and so know from (4) the isolations will be dominated by:

$$\left(2d_E\right)^{3^{n+1}} m_E^{2^n} d_E^8 + \left(\left(2d_E\right)^{3^{n+1}} m_E^{2^n}\right)^7 d_E^7 l_E^3. \tag{15}$$

By combining (15) with (13) and (14) we are now in a position to describe the overall complexities of our algorithms. Note that (15) will be a large overestimation for the 1-layered case.

Theorem 6.
The complexity for computing a variety sub-CAD using Collins' algorithm is dominated by:

$$2^{12} \left(2^2 d_A^2\right)^{\mathbf{4^{n+3}}} \left(\frac{m_E(2d_E + 2m_A - 1)}{2}\right)^{\mathbf{2^{n+5}}} l_A^3 d_A^{12} +$$

$$\left(2d_E\right)^{3^{n+1}} m_E^{2^n} d_E^8 + \left(\left(2d_E\right)^{3^{n+1}} m_E^{2^n}\right)^7 d_E^7 l_E^3. \tag{16}$$

The complexity for computing a 1-layered variety sub-CAD is dominated by:

$$2^{12} \left(2^2 d_A^2\right)^{\mathbf{3^{n+3}}} \left(\frac{m_E(2d_E + 2m_A - 1)}{2}\right)^{\mathbf{2^{n+3}}} l_A^3 d_A^{12} +$$

$$\left(2d_E\right)^{3^{n+1}} m_E^{2^n} d_E^8 + \left(\left(2d_E\right)^{3^{n+1}} m_E^{2^n}\right)^7 d_E^7 l_E^3. \tag{17}$$

In Theorem 6 we have emboldened the exponents to highlight the difference between (16) and (17). To help visualise the comparison we have plotted the double logarithm of the complexities against n in Figure 5 (for some specific parameter values). The diagram shows the drop in the constant in the exponents of (16) and (17), whilst the scaling factor of the exponent remains the same between variety and non-variety versions of each algorithm.

5. Examples and Implementation

We provide some case studies showing the benefit of our new algorithms, which have all been implemented in the Maple package `ProjectionCAD` [24, 25, 45].

We also compare to some competing CAD implementations: the CAD procedures in Maple's `RegularChains` Library, the command line program Qepcad [8], and the algorithm in Mathematica [42] (which produces a cylindrical algebraic formula). We tested two `RegularChains` routines: the one following [16] (distributed with Maple), and the one following [15] which can make use of equational

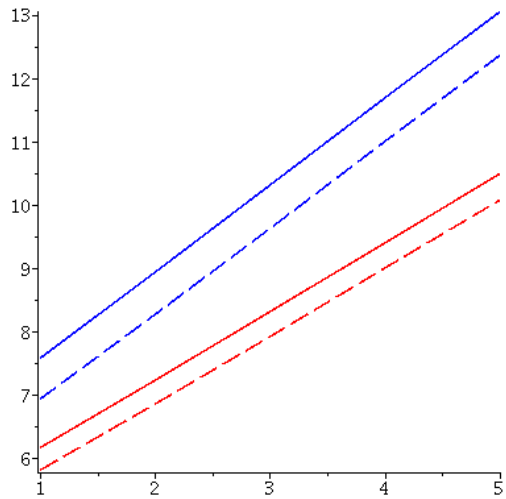

FIGURE 5. A plot of n (horizontal axis) against the double logarithm of the complexities of algorithms (vertical axis). The complexities were evaluated with parameter choices $d_A = 3$, $d_E = 2$, $m_A = 3$, $m_E = 1$, $m_{A\setminus E} = 2$, $l_A = 2$, $l_E = 2$. From top to bottom: CAD, variety sub-CAD, 1-layered sub-CAD and 1-layered variety sub-CAD.

constraints. With QEPCAD we ran it first on with its default settings (implementing [33]) and also when an equational constraint is designated (where it follows [34]).

QEPCAD also has the option `measure-zero-error` which produces a CAD with only the full-dimensional cells of the free variable space guaranteed to satisfy the invariance condition (see [8]). Although this is a CAD rather than a sub-CAD of the free variable space it is only sufficient to ensure the full dimensional solutions are correct, like a 1-layered sub-CAD, meaning the solutions have an *error of measure zero in the free-variable space*. For the three case studies below this command was not of use because the problems were unquantified and each have an equational constraint, meaning their formulae can only be satisfied on cells of less than full-dimension (in the free variable space). So although we could produce an output from QEPCAD using `measure-zero-error`, the only cells guaranteed to be correct are not of interest. This contrasts with a 1-layered variety sub-CAD which does provide valid solutions. Here the layer is with respect to the variety, not free variable space, and so the solutions provided have an *error of measure zero in the solution space*.

Experiments were run on a Linux desktop (3.1GHz Intel processor, 8.0Gb total memory). Tests for MATHEMATICA used V9 and for QEPCAD used QEPCAD-B 1.69 with the options `+N500000000` and `+L200000` (initialisation times included). The tests in MAPLE used the development version (similar to MAPLE 18) in command line interface, using the development version of the `RegularChains` Library[1].

5.1. Example: Making use of a 1-layered variety sub-CAD

Assume variable ordering $x \succ y \succ z$ and consider the following formula involving 3 random polynomials of degree 2 (generated using MAPLE's `randpoly` function) which are plotted in Figure 6:

$$\begin{aligned}\Phi := & -50xy + 56yz + 41z^2 + 67x - 55y - 21 = 0 \\ \wedge \quad & 36xy + 76xz - 58yz + 69z^2 + 75y + 27 > 0 \\ \wedge \quad & -55x^2 + 10xy - 88x + 80y + z - 39 > 0.\end{aligned}$$

We wish to describe the regions of $\mathbb{R}^3$ in which Φ is satisfied. Figure 6 shows there are multiple intersections between the two non-equational constraints away from the variety defined by the equational

[1]Available from `www.regularchains.org`.

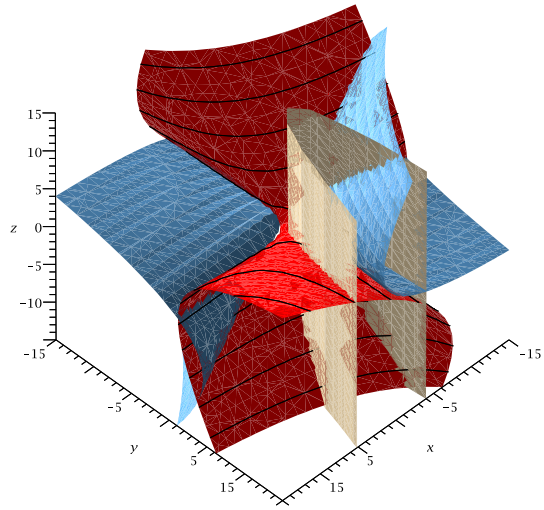

FIGURE 6. Intersection of the three surfaces from Section 5.1. The red surface (the darkest in black and white) is the equational constraint.

constraint. This suggests that V-sub-CAD could be beneficial. Further, if only generic solution sets are of interest then a 1-layered variety sub-CAD could offer further savings.

To show the relative benefits of the individual theories we solved the problem by constructing: a full sign-invariant CAD for the three polynomials (following [33]); a CAD invariant with respect to the equational constraint (following [34]); a variety sub-CAD (following Algorithm 1 implementing sub-procedures from [34]) and layered variety sub-CADs (following Algorithm 5 implementing sub-procedures from [34]).

Full sign-invariant CAD. 17,047 cells, 178.277 seconds.
CAD invariant with respect to EC. 1315 cells, 11.520 seconds.
Variety sub-CAD. 422 cells, 10.723 seconds.
2-Layered Variety Sub-CAD. 348 cells, 7.149 seconds.
1-Layered Variety Sub-CAD. 138 cells, 0.475 seconds.

We see that making use of the equational constraint in the projection stage dramatically reduces the computation involved. A variety sub-CAD further reduces the size of the output which will lead to time savings on any future work. The variety sub-CAD is sufficient to describe exactly where Φ is true, but if we are only concerned with the generic solution sets then a 1-layered variety sub-CAD can be used to achieve further time-savings. Note that this means an output of over 17,000 cells can be replaced with one of only 138. If solutions of lower dimension are also needed then the 2-layered variety sub-CAD (348 cells) or the complete variety sub-CAD (422 cells) also offer great savings.

We now consider competing CAD implementations:

Maple: Algorithm from [16]**.** 9841 cells, 112.460 seconds.
Maple: Algorithm from [15]**.** 559 cells, 1.999 seconds.
Qepcad. 17,047 cells, 385.679 seconds.
Qepcad EC. 5271 cells, 26.614 seconds.
Mathematica. 0.533 seconds

We see that those algorithms taking advantage of the equational constraint offer smaller quicker CADs. QEPCAD does comparatively worse here (perhaps due to the lack of improved lifting described in [5]).

If we did use QEPCAD with its `measure-zero-error` option here it would return (in under 5 seconds) a CAD with 822 cells, but Φ would not be satisfied on any of them (as expected) and hence QEPCAD gives an equivalent quantifier free formula `False`. For the QEPCAD outputs above Φ is only valid on a small fraction of cells (290/17047 and 106/5271). The false cells that are constructed (and on which Φ is evaluated) will make a significant contribution to the computation time. We can evaluate Φ on all the cells produced in the 1-LV-sub-CAD almost instantly to find that there are 36 of 138 cells on which Φ is satisfied (the equation is by definition satisfied for all of them but the truth of the other constraints varies).

MATHEMATICA does not produce CAD cells, and so we cannot compare cell counts. Instead it produces a cylindrical formula, from which the solutions can be derived. This is done very quickly, (probably due to the symbolic-numeric techniques discussed in [41]).

Consider a general problem of the form $f = 0 \wedge \Psi(g_i)$ where $f = 0$ defines a variety of *real* co-dimension 1 and Ψ is a quantifier-free formula involving only $f = 0$ and strict inequalities. We expect that a 1-layered variety sub-CAD will usually be sufficient to describe the generic solutions. We may find that there are no solutions of full dimension on the variety, and can use the recursive approach for layered CAD to incrementally build extra layers as required.

5.2. Example: Making use of 1-layered variety sub-TTICAD

We now consider a problem suitable for both our new sub-CAD approaches and TTICAD. Define the following quantifier free formulae:

$$\begin{aligned}\varphi_1 &:= x^2 + y^2 + z^2 = 1 \wedge xy + yz + zx < 1 \wedge x^3 - y^3 - z^3 < 0,\\ \varphi_2 &:= (x-1)^2 + (y-1)^2 + (z-1)^2 = 1 \wedge (x-1)(y-1) + (y-1)(z-1)\\ &\qquad + (z-1)(x-1) < 1 \wedge (x-1)^3 - (y-1)^3 - (z-1)^3 < 0.\end{aligned}$$

The surfaces defined by the polynomials in φ_1 are shown in Figure 7, while those in φ_2 are the same but shifted. Assume the variable ordering $x \succ y \succ z$ and consider the problem of finding all regions of $\mathbb{R}^3$ satisfying $\Phi := \varphi_1 \vee \varphi_2$. It would be naïve to tackle this problem with a sign-invariant CAD for the 6 polynomials in Φ, (the default CAD in MAPLE 16 could not produce one when left overnight while QEPCAD gives a "prime list exhausted" error (a memory constraint) after two hours).

Instead we should make use of the equations present in the formulae. We can do this with QEPCAD by declaring the implicit equational constraint (the product of the two equations) to give a CAD with 6165 cells in 35,304 seconds (around 10 hours). MATHEMATICA can produce a cylindrical formula in 1.805 seconds. This problem is well suited for TTICAD and applying the `ProjectionCAD` implementation on the two formulae φ_1 and φ_2 produces 4861 cells in 170.515 seconds. A substantially lower cell count than QEPCAD is achieved because the TTICAD projection set is smaller than the one using the implicit equational constraint (see [4, 5]).

We now consider how the TTICAD can be improved upon using the new theory. Each formula φ_i contains an equational constraint and so the formula Φ is only true on the variety defined by their product. Suppose further that we only want to obtain the generic solutions. Then we may apply the TTICAD operator to φ_1 and φ_2 and construct a 1-layered sub-CAD of $\mathbb{R}^2$ with respect to this projection set. This takes 0.947 seconds and produces 249 cells in $\mathbb{R}^2$. We then lift with respect to both of the equational constraints onto the variety defined by their product. This takes a further 1.191 seconds and produces 528 2-dimensional cells on the 2-dimensional variety in $\mathbb{R}^3$. So the 1-layered variety sub-TTICAD saves 88% of the cells and 99% of the computation time of the TTICAD.

The 1-layered variety sub-TTICAD obtains all the cells of full dimension (with respect to the variety) on which Φ is true, and so is sufficient up to a set of measure zero. If solutions of lower dimension are needed then the theory in this paper allows for a 2-layered variety sub-TTICAD or the full variety sub-TTICAD which contain 1514 cells and 1976 cells, respectively. The former takes under a minute and the latter just less than three to compute. So we see that the variety sub-TTICAD takes about the same time as a TTICAD, but gives a cell saving, while the layered variety sub-TTICADs

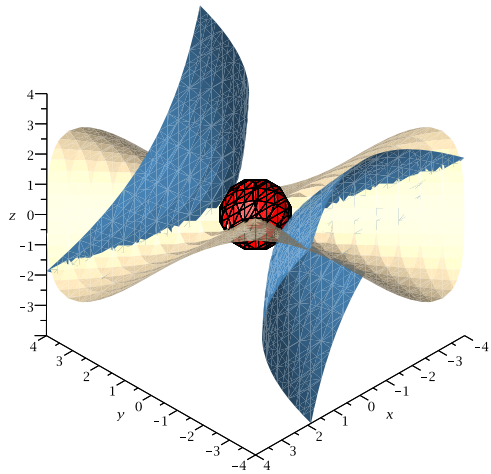

FIGURE 7. Intersection of the surfaces from φ_1 in Section 5.2 – the sphere defined by the equational constraint is in red (the darkest surface in black and white).

also offer time savings. Note that over half the cells in the complete TTICAD do not lie on either of the varieties defined by the equational constraints.

To magnify the issues consider a third formula (a different shift of the original surfaces)

$$\varphi_3 := (x+1)^2 + (y+1)^2 + (z+1)^2 = 1 \wedge (x+1)(y+1) + (y+1)(z+1) + (z+1)(x+1) < 1 \wedge (x+1)^3 - (y+1)^3 - (z+1)^3 < 0,$$

and a new overall formula, $\Phi^* := \varphi_1 \vee \varphi_2 \vee \varphi_3$.

Given the above results for Φ we do not attempt to solve this problem without utilising the equational constraints. This time to build a 1-layered variety sub-TTICAD takes 5.003 seconds and produces 1104 cells. The full TTICAD takes 432.210 seconds and produces 10063 cells; around 10 times as many. Although the TTICAD will contain all valid cells for Φ, the 1-layered variety sub-TTICAD will provide descriptions of the generic families of solutions. If solutions of lower dimensions are needed then a 2-layered variety sub-TTICAD or the complete variety sub-TTICAD would both be preferable to the full TTICAD. The former contains 3166 cells and took 145.898 seconds, and the latter contains 4130 cells and took 429.083 seconds.

5.3. Example: A "Piano Movers" problem

An application of CAD of great interest is motion planning. Given a semi-algebraic object, an initial and desired position, and semi-algebraic obstacles, a CAD can be constructed of the valid configuration space of the object. The connectivity of this space can then be used to determine if a feasible path from the initial position of the object to the desired endpoint is possible [38].

A well-studied problem in this area is the movement of a ladder through a right-angled corridor, as proposed in [20]. The problem consists of a ladder of length 3 inside a right-angled corridor of width 1 with the aim of moving from from position 1 to position 2 in Figure 8.

Although in 2-dimensional real space, the problem lies in a 4-dimensional configuration space (as usually the problem is described using 4 variables specifying the endpoints of the ladder) making it far more difficult to describe using CAD. The original formulation given in [20] proves particularly difficult but recently in [44] a reformulation beneficial to CAD was given. A CAD of the configuration space was built using QEPCAD with 285,419 cells in around 5 hours. This used the equational constraint

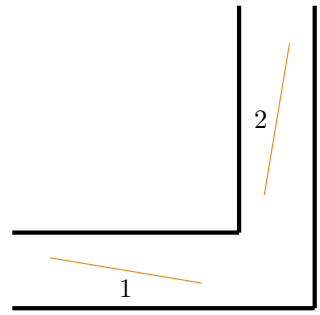

FIGURE 8. The piano movers problem considered in [20]

[34] and partial CAD techniques [18] (without these it increases to 1,691,473 cells and over 24 hours computation time).

A representation of the 2-dimensional CAD produced on route to the full 4-dimensional CAD is given in Figure 9. MATHEMATICA can produce a cylindrical formula in 558.721 seconds, but for this application such a formula is not sufficient to deduce paths (since that will require knowledge of cell adjacencies and thus the boundary cells which can not always be inferred from the formula).

In the formulation the length of the ladder is an explicit equational constraint, and so the problem is suited to treatment with a V-sub-CAD. Indeed, we can see from Figure 9 that there is a great deal of information computed which is of no use in describing the suitable paths. Ideally we would restrict to a CAD or sub-CAD of just the corridor highlighted.

Within this three-dimensional variety (embedded within $\mathbb{R}^4$) the important cells are those that are three-dimensional, since cells of lesser dimension correspond to physically infeasible situations (i.e. one dimensional subspaces of $\mathbb{R}^2$). Therefore a LV-sub-CAD would give an overview of the problem.

Constructing the 1-layered sub-CAD of $\mathbb{R}^3$ produces 64,764 cells in around 124.22 seconds, and lifting to the variety takes a further 196.672 seconds, producing 101,924 cells, so offering substantial savings. Since this example has several constraints which do not contain all variables it is likely that partial CAD techniques (not yet implemented in our MAPLE package) would offer further savings. Adjacency information will have to be computed for a full solution to the motion planning problem (which may require a 2-layered variety sub-CAD).

6. Conclusions and further work

We have formalised the idea of a cylindrical algebraic sub-decomposition. Whilst a simple idea, it can be hugely powerful and the examples presented show that massive cell reductions are possible. In some cases time reductions are also available, and since most problems using CAD will require some further computation on the cells (such as polynomial evaluation) more time savings will follow and applications involving complicated calculation on the cells will benefit even more. For example, the calculation of adjacency information for use in motion planning or the evaluation of multi-valued functions at (possibly algebraic) sample points for branch cut analysis.

We provided examples of sub-CADs in the literature, along with two new approaches (and algorithms to produce them): variety sub-CADs and layered sub-CADs. We find that their individual savings may be magnified by combining them with each other and the recent theory of truth table invariance. The savings are large enough to tackle problems previously infeasible.

There is great scope for future work, with some important questions as follows:

- Can we identify further classes of problems where the various types of sub-CAD are sufficient?
- What is the best way to build V-sub-CADs for lower dimensional varieties?
- How can we best adapt existing techniques (such as partial CAD) to output sub-CADs?

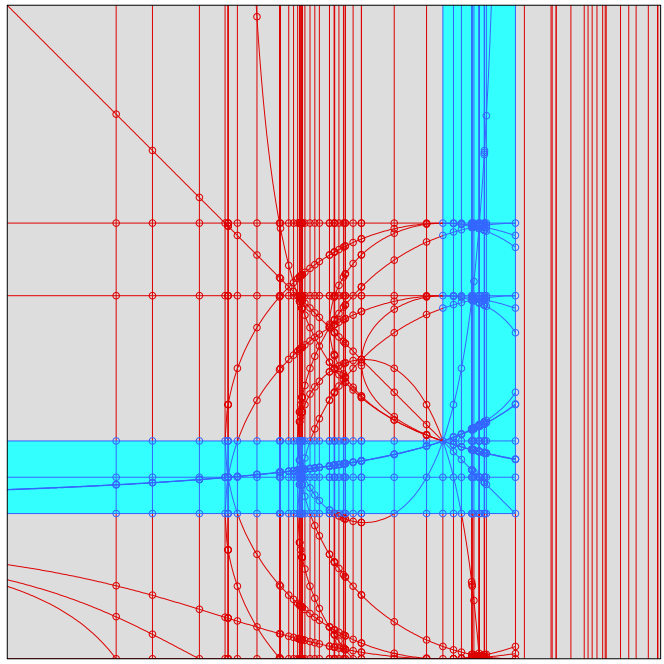

FIGURE 9. A representation of the 2-dimensional induced CAD produced for the new formulation of the piano movers problem in [44].

- Can we develop heuristics (or adapt existing ones [23, 6]) for when to use different sub-CAD approaches?
- Can we keep track of where cells arise when constructing a variety sub-TTICAD so that over each cell we lift only to the varieties for relevant ϕ_i? This can be thought of as an analogue of partial CAD for TTICAD. This may alter the output significantly (it may not be a sub-CAD of a CAD that can be constructed by current technology) but could allow for even smaller output for suitable problems.
- Can we parallelise the algorithms? The idea mentioned in [32] of lifting over sets of cells independently could be generalised. (Preprocessing CAD problems to allow for parallelisation was discussed in [30] but this involved only the boolean logic of the problem).

There are also interesting questions around which properties of CADs transfer over to their sub-CADs. For example, any existing adjacency algorithms require CADs to be particularly 'well-behaved', and it may be possible to avoid problematic cells through sub-CADs, extending the use of such algorithms. Also, well-orientedness conditions for CAD algorithms [33, 7] may be failed for the CAD but not the sub-CADs. These ideas need further investigation.

An overarching aim is to develop a general CAD framework to identify when applying each technique is appropriate, automatically combining appropriate methods when possible and making choices automatically when required based on heuristic information. This would identify, for a given problem, an efficient way to produce a ϕ-sufficient sub-CAD and thus describe the solution set.

Acknowledgements

This work was supported by the EPSRC grant: EP/J003247/1. The authors would also like to thank Professor Gregory Sankaran for his thoughts and feedback on the topic, and Professor Scott McCallum for many stimulating conversations on TTICAD. Finally, they would like to thank the anonymous referees for helpful comments which improved the paper.

References

[1] D. Arnon, G.E. Collins, and S. McCallum. Cylindrical algebraic decomposition I: The basic algorithm. *SIAM Journal of Computing*, 13:865–877, 1984.

[2] J. Backelin. Square multiples n give infinitely many cyclic n-roots. Matematiska Institutionen Reports Series, Stockholms Universitet, 1989.

[3] R. Bradford and J.H. Davenport. Towards better simplification of elementary functions. In *Proc. ISSAC '02*, pages 16–22. ACM, 2002.

[4] R. Bradford, J.H. Davenport, M. England, S. McCallum, and D. Wilson. Cylindrical algebraic decompositions for boolean combinations. In *Proc. ISSAC '13*, pages 125–132. ACM, 2013.

[5] R. Bradford, J.H. Davenport, M. England, S. McCallum, and D. Wilson. Truth table invariant cylindrical algebraic decomposition. *Submitted*. Preprint: `http://opus.bath.ac.uk/38146/`, 2014.

[6] R. Bradford, J.H. Davenport, M. England, and D. Wilson. Optimising problem formulations for cylindrical algebraic decomposition. In J. Carette, D. Aspinall, C. Lange, P. Sojka, and W. Windsteiger, editors, *Intelligent Computer Mathematics*, (LNCS 7961), pages 19–34. Springer Berlin Heidelberg, 2013.

[7] C.W. Brown. Improved projection for cylindrical algebraic decomposition. *Journal of Symbolic Computation*, 32(5):447–465, 2001.

[8] C.W. Brown. An overview of QEPCAD B: A program for computing with semi-algebraic sets using CADs. *SIGSAM Bulletin*, 37(4):97–108, ACM, 2003.

[9] C.W. Brown. The McCallum projection, lifting, and order-invariance. Technical report, U.S. Naval Academy, Computer Science Department, 2005.

[10] C.W. Brown. Constructing a single open cell in a cylindrical algebraic decomposition. In *Proc. ISSAC '13*, pages 133–140. ACM, 2013.

[11] C.W. Brown and J.H. Davenport. The complexity of quantifier elimination and cylindrical algebraic decomposition. In *Proc. ISSAC '07*, pages 54–60. ACM, 2007.

[12] C.W. Brown, M. El Kahoui, D. Novotni, and A. Weber. Algorithmic methods for investigating equilibria in epidemic modelling. *Journal of Symbolic Computation*, 41:1157–1173, 2006.

[13] C.W. Brown and S. McCallum. On using bi-equational constraints in CAD construction. In *Proc. ISSAC '05*, pages 76–83. ACM, 2005.

[14] M.A. Burr. Applications of continuous amortization to bisection-based root isolation. *Preprint*: `http://arxiv.org/abs/1309.5991`, 2013.

[15] C. Chen and M. Moreno Maza. An incremental algorithm for computing cylindrical algebraic decompositions. *Proc. ASCM '12*, (to appear, Springer). Preprint: `arXiv:1210.5543`, 2012.

[16] C. Chen, M. Moreno Maza, B. Xia, and L. Yang. Computing cylindrical algebraic decomposition via triangular decomposition. In *Proc. ISSAC '09*, pages 95–102. ACM, 2009.

[17] G.E. Collins. Quantifier elimination for real closed fields by cylindrical algebraic decomposition. In *Proc. 2nd GI Conference on Automata Theory and Formal Languages*, pages 134–183. Springer-Verlag, 1975.

[18] G.E. Collins and H. Hong. Partial cylindrical algebraic decomposition for quantifier elimination. *Journal of Symbolic Computation*, 12:299–328, 1991.

[19] J.H. Davenport. Computer algebra for cylindrical algebraic decomposition. Technical Report TRITA-NA-8511, NADA KTH Stockholm. Reissued as Bath Computer Science Technical report 88-10. Available at `http://staff.bath.ac.uk/masjhd/TRITA.pdf`, 1985.

[20] J.H. Davenport. A "Piano-Movers" Problem. *SIGSAM Bulletin*, 20(1-2):15–17, 1986.

[21] J.H. Davenport, R. Bradford, M. England, and D. Wilson. Program verification in the presence of complex numbers, functions with branch cuts etc. In *Proc. SYNASC '12*, pages 83–88. IEEE, 2012.

[22] J.H. Davenport and J. Heintz. Real quantifier elimination is doubly exponential. *Journal of Symbolic Computation*, 5(1-2):29–35, 1988.

[23] A. Dolzmann, A. Seidl, and T. Sturm. Efficient projection orders for CAD. In *Proc. ISSAC '04*, pages 111–118. ACM, 2004.

[24] M. England. An implementation of CAD in Maple utilising McCallum projection. Department of Computer Science Technical Report series 2013-02, University of Bath.
Available at `http://opus.bath.ac.uk/33180/`, 2013.

[25] M. England. An implementation of CAD in Maple utilising problem formulation, equational constraints and truth-table invariance. Department of Computer Science Technical Report series 2013-04, University of Bath. Available at `http://opus.bath.ac.uk/35636/`, 2013.

[26] M. England, R. Bradford, J.H. Davenport, and D. Wilson. Understanding branch cuts of expressions. In J. Carette, D. Aspinall, C. Lange, P. Sojka, and W. Windsteiger, editors, *Intelligent Computer Mathematics*, (LNCS 7961), pages 136–151. Springer Berlin Heidelberg, 2013.

[27] I.A. Fotiou, P.A. Parrilo, and M. Morari. Nonlinear parametric optimization using cylindrical algebraic decomposition. In *Decision and Control, 2005 European Control Conference. CDC-ECC '05.*, pages 3735–3740, 2005.

[28] H. Hong. An improvement of the projection operator in cylindrical algebraic decomposition. In *Proc. ISSAC '90*, pages 261–264. ACM, 1990.

[29] H. Iwane, H. Yanami, H. Anai, and K. Yokoyama. An effective implementation of a symbolic-numeric cylindrical algebraic decomposition for quantifier elimination. In *Proc. SNC '09*, pages 55–64, 2009.

[30] H.K. Malladi and A. Dukkipati. A preprocessor based on clause normal forms and virtual substitutions to parallelize cylindrical algebraic decomposition. *Preprint:* `http://arxiv.org/abs/1112.5352v3`, 2013.

[31] S. McCallum. Solving polynomial strict inequalities using cylindrical algebraic decomposition. *The Computer Journal*, 36(5):432–438, 1993.

[32] S. McCallum. A computer algebra approach to path finding in the plane. In J. Harland, editor, *Proceedings of Computing: The Australasian Theory Symposium (CATS)*, pages 44–50, 1997.

[33] S. McCallum. An improved projection operation for cylindrical algebraic decomposition. In B. Caviness and J. Johnson, editors, *Quantifier Elimination and Cylindrical Algebraic Decomposition*, Texts & Monographs in Symbolic Computation, pages 242–268. Springer-Verlag, 1998.

[34] S. McCallum. On projection in CAD-based quantifier elimination with equational constraint. In *Proc. ISSAC '99*, pages 145–149. ACM, 1999.

[35] S. McCallum. On propagation of equational constraints in CAD-based quantifier elimination. In *Proc. ISSAC '01*, pages 223–231. ACM, 2001.

[36] L.C. Paulson. Metitarski: Past and future. In L. Beringer and A. Felty, editors, *Interactive Theorem Proving*, (LNCS 7406), pages 1–10. Springer, 2012.

[37] N. Phisanbut, R.J. Bradford, and J.H. Davenport. Geometry of branch cuts. *ACM Communications in Computer Algebra*, 44(3):132–135, 2010.

[38] J.T. Schwartz and M. Sharir. On the "Piano-Movers" Problem: II. General techniques for computing topological properties of real algebraic manifolds. *Adv. Appl. Math.*, 4:298–351, 1983.

[39] A. Seidl and T. Sturm. A generic projection operator for partial cylindrical algebraic decomposition. In *Proc. ISSAC '03*, pages 240–247. ACM, 2003.

[40] A. Strzeboński. Solving systems of strict polynomial inequalities. *Journal of Symbolic Computation*, 29(3):471–480, 2000.

[41] A. Strzeboński. Cylindrical algebraic decomposition using validated numerics. *Journal of Symbolic Computation*, 41(9):1021–1038, 2006.

[42] A. Strzeboński. Computation with semialgebraic sets represented by cylindrical algebraic formulas. In *Proc. ISSAC '10*, pages 61–68. ACM, 2010.

[43] A. Strzeboński. Solving polynomial systems over semialgebraic sets represented by cylindrical algebraic formulas. In *Proc. ISSAC '12*, pages 335–342. ACM, 2012.

[44] D. Wilson, J.H. Davenport, M. England, and R. Bradford. A "piano movers" problem reformulated. In Proc. SYNASC '13, IEEE, 2013.

[45] D. Wilson and M. England. Layered cylindrical algebraic decomposition. Department of Computer Science Technical Report series 2013-05, University of Bath. Available at `http://opus.bath.ac.uk/36712/`, 2013.

D. J. Wilson, R. J. Bradford, J. H. Davenport and M. England
Department of Computer Science, University of Bath, Bath, BA2 7AY, England
e-mail: {D.J.Wilson, R.J.Bradford, J.H.Davenport, M.England}@bath.ac.uk